

Cloud Computing Special Interest Group

Cloud Computing for the UK Research Community
Workshop

28-29 December 2013

Philip Kershaw, STFC Rutherford Appleton Laboratory

Overview

- Who is the Cloud SIG?
- Why and why now?
- Workshop structure
- Desired outcomes

Cloud SIG Committee Members

- David Blundell, 100 Percent IT
- John Chapman, JANET
- Neil Chue-Hong, Software Sustainability Institute
- David Colling, Imperial College
- David Fergusson, Crick Institute
- Roger Jones, CERN
- Philip Kershaw, STFC Rutherford Appleton Laboratory (chair)
- David Wallom, Oxford e-Research Centre
- Jeremy Yates, UCL

Origins

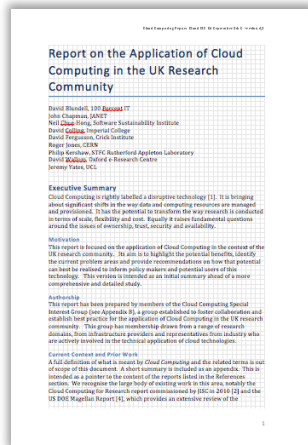
- Formed as a follow up from e-Infrastructure Project Directors Group meeting, July this year
 - Identified key areas where it was felt there could be benefit from greater co-ordination
 - e.g. cloud computing, identity management and access control . . .
- The SIG is an independent group
- A second incarnation
 - Original SIG co-ordinated by Neil Chue Hong (SSI) and Jeremy Cohen (Imperial College)

Why?

- Too many SIGs already?
- The effectiveness of a SIG:
 - the need to actively create, co-ordinate and enable
- We seek to have a direct link with and support projects which are already underway in this space
- There is much existing work in this area already
 - Reports: Magellan (DoE 2011), Cloud Computing for Research (JISC, 2010), ...
- But it is a fast moving area with new challenges


Challenges and Opportunities

- The challenges of *Big Data* are impacting on a wider set of communities within the research domain.
 - Cloud computing provides *one* solution to tackle these challenges the ability to
 - to scale out compute and storage
 - to bring users to data through the provision of hosted processing and analysis environments at data centres
- Opportunities provided through recent investment from UK government in e-Infrastructure
- Increased maturity of software for private cloud provision (e.g. OpenStack and vCloud)
- Increased availability of pre-packaged Virtual Appliances for scientific computing (e.g. Galaxy CloudMan)



Cloud SIG Executive Summary

Issues to Address



Gartner Hype Cycle

- Standardisation of interfaces
 - multiple competing standards, a lack of standards, incomplete standards or existing standards with insufficient uptake
- SLAs and charging models
 - need for greater understanding, education for users
 - Difference between commercial and research funding models
 - Can public providers make custom SLA arrangements for research sector?
- Danger of fragmentation within research community
 - Bespoke solutions developed for generic problems
- ...

Terms of Reference

The SIG has a number of goals as follows:

- To establish best practice for the application of cloud computing in the *UK research community*
- Share this best practice and provide recommendations where needed
- *Execute* these recommendations
- Focus work around *existing projects* in the research community to ensure a direct connection to practical application of cloud technologies.
- Co-ordinate efforts and foster collaboration across research communities seeking to avoid domain-specific solutions

Workshop Structure

- Structured around three themes
 - Public cloud
 - How can research users best exploit public cloud resources?
 - Give provider and consumer perspectives
 - Private cloud
 - A number of organisations from the research sector are considering or in the process of rolling out a private cloud
 - How can a cloud best integrate with existing infrastructure and functions? - data centre, HPC, hosted processing and analysis environments
 - Cloud federation and brokering
 - What are the opportunities that are offered from procuring resources from multiple providers?
 - What are the technical and policy-related challenges?
 - Can brokers support the user community to get the best service?

Outcomes from Workshop

- Identify a set of themes, common problem areas
- Select priority areas
- Use pilot projects as practical vehicles to
 - explore these and find solutions or
 - provide momentum for solutions
- These should be existing projects already running or planned in the research community
- Provide a single voice to lobby for change
- (to explore further in final session tomorrow)

Summary

- Public cloud
 - Ability to massively scale-out
 - Immediate availability
 - Trust
 - WAN b/w – data in / data out
 - Service models and SLA that fits the research community's needs
- Federation and brokering
 - Technical:
 - opportunity to select from a range of providers with a common interface
 - technical challenges, evolving capability
 - Policy: possibility of collective bargaining
- Private cloud
 - Capex vs. opex with open source vs. proprietary
 - Complexity in deployment, available features, stability
 - Licensing
 - Trajectory of s/w vendor be it open source or proprietary

Next Steps

- SMART objectives
 - Specific, Measurable, achievable, realistic, time-bound
- What common themes do we have across the community?
 - Private cloud expertise
 - Hybrid and federation: APIs, virtual appliance portability
 - Charging models, SLAs
- Training
 - Help users - Making cloud friendly tools
- Policy-based?
- Technical?
 - Document tech best practice to eliminate misunderstanding about what cloud can and can't do
- What assets do we have as SIG?
 - collective input and co-ordination
 - Exchange of expertise and best practice
- Know who your users are
 - IaaS and PaaS, SaaS are different users: admins, developers/researchers, end users
- SIG put proposals BUFDG (British Universities Finance Directors Group)
 - Best practice about financing cloud computing for research